

XCELL

THE NEWSLETTER FOR XILINX PROGRAMMABLE LOGIC USERS

Issue 8

Fourth Quarter 1992

This issue is mainly dedicated to hardware considerations. Most important is the doubling of FPGA performance in the new XC3100 family, available now. More accurate figures for dynamic power consumption are available, and we now guarantee values for the shortest, or best case, delays. Two tutorial articles deal with transmission line effects on PC-boards, and with crystal oscillators. We also give a brief description of the major software release that goes out in Nov./Dec.

Peter Alfke, editor

Table of Contents

Component Availability	2
Software	
New Software Release	4
APR Speed on diff. Platforms	5
Victory for JTAG	5
Hardware	
3100 Switching Characteristics	3
Legal Protection of Bit Streams	5
EPLD Architecture	6
Dynamic Power Consumption	7
Volatility	8
Battery Backup	8
Minimum Logic Delays Guaranteed	9
Master and Slave Configure Together	9
128K-Bit Serial PROM	9
XC4000 Data Book Reprint	9
PC-Board Design Hints	10
Crystal Oscillators	12
Simple RC Oscillator	14
Metastability Equations	14
Shorter Power-On Delay	15
255-MHz Frequency Synthesizer	16

Faster than 200 MHz

In early November 1992, Xilinx will formally announce the new XC3100 family. It offers familiar functionality and packages, is 100% software and pin-out compatible with the XC3000 family, but can run twice as fast. Samples are available now.

The five members of this family, XC3120, XC3130, XC3142, XC3164, and XC3190 each come in three speed grades, -5, -4, and -3. The nomenclature is similar to that of the XC4000 family; the speed number is a rounded-up version of the CLB-propagation delay, in nanoseconds.

The slowest speed grade, XC3100-5, is 25% faster than the XC3000-125; the middle speed grade is another 25% faster, the equivalent of a non-existing XC3000-190. The top speed grade, XC3100-3, is another 25% faster and has a T_{IO} of 2.7 ns and can toggle (worst-case, guaranteed over temperature and voltage,) at >230 MHz. To prove the point, we built a 255-MHz frequency synthesizer running inside the XC3130-3. (See back page).

This revolutionary increase in performance is the result of several evolutionary developments.

- The XC3100 family is manufactured using an aggressive 0.8 μ process optimized for LCA logic and routing

- The circuit design has been carefully tuned to achieve shorter delays. Special attention has been

paid to the pass transistors in the interconnect structure. Their resistance has been cut in half, resulting in significantly faster interconnect paths. The logic inside the CLBs was also made faster by carefully matching the transistors to the capacitive load.

Advanced processing and improved circuit design result in a doubling of performance, compared to the present XC3000-125 family. This again demonstrates the versatility of the industry-standard CMOS-SRAM process. As technology advanced from 1.2 μ to 0.8 μ minimum feature size, Xilinx took advantage of the improved transistor parameters, reducing the stray capacitance and on-resistance of the pass transistors, slashing the RC interconnect delay by a factor 3.

This opens up new applications, where FPGAs previously could not compete. Microprocessors have become faster; 50-MHz clock rates are now common. The new XC3100 family meets this challenge; designers can now incorporate programmable logic in demanding high-speed systems

Converting existing XC3000 designs to the new XC3100 family is a cinch; just plug in the new part. It is 100% software, hardware and package pin-out compatible. Several users have already made the switch, and have confirmed our characterization data. P.A.

LCA Component Availability (November 1992)

	44 PIN	68 PIN	84 PIN		100 PIN			120 PIN		132 PIN		156 PIN		160 PIN		164 PIN		175 PIN		191 PIN	196 PIN	208 PIN
	PLASTIC PLCC	PLASTIC PLCC	CERAMIC PGA	PLASTIC PLCC	CERAMIC PGA	PLASTIC PQFP	CERAMIC COFP	PLASTIC TOFP	CERAMIC PGA	PLASTIC PGA	CERAMIC COFP	CERAMIC PGA	PLASTIC PQFP	CERAMIC COFP	BRAZED CERAMIC COFP	TOP BRAZED CERAMIC COFP	PLASTIC PGA	CERAMIC PGA	CERAMIC PGA	TOP BRAZED CERAMIC COFP	PLASTIC PQFP	
	PC44	PC68	PG68	PC84	PG84	PC100	CD100	TQ100	PG120	PP132	PG132	PG156	PC160	CD164	CB164	PP175	PG175	PG185	GB196	PQ208		
XC2064	-50	CI	CI	CIM																		
	-70	CI	CI	CIM																		
	-100	C	C	C																		
XC2018	-33					MB																
	-50	CI	CI			CI	CIMB															
	-70	CI	CI			CI	CIMB															
	-100	C	C			C	C								C							
XC3020	-70		CI			CI	CIMB	CI	CIMB													
	-100		CI			CI	CIM	C	C													
	-125		C			C	C	C	C													
XC3030	-70	CI	CI			CI	CIM	C														
	-100	C	CI			CI	CIM	C							C							
	-125	C	C			C	C	C	C						C							
XC3042	-70					CI	CIMB	CI	CMB						CI	CI						
	-100					CI	CIM	CI	CM	C					CI	CIM						
	-125					CI	CI	CI	C	C					CI	CI						
XC3064	-70					C									CI	CIM	C					
	-100					CI									CI	CIM	C					
	-125					C									C	C	C					
XC3090	-70					CI									CI	CMB	MB	CI	CIMB	CI		
	-100					CI									CIM	C	M	CI	CIM	CI		
	-125					C									C		C	C	C	C		
XC3120	-5	CI				CI	CI	CI	CI													
	-4	CI				CI	CI	CI	CI													
	-3	CI				C	C	C	C													
XC3130	-5	CI	CI			CI	CI	CI	CI													
	-4	CI	CI			CI	CI	CI	CI													
	-3	G	G			C	C	C	C													
XC3142	-5					CI	CI	CI	CI						CI	CI						
	-4					CI	CI	CI	CI						CI	CI						
	-3					C	C	C	C						C	C						
XC3164	-5					CI									CI	CI	CI					
	-4					CI									CI	CI	CI					
	-3					C									C	C	C					
XC3190	-5					CI									CI		CI	CI	CI	CI		
	-4					CI									CI		CI	CI	CI	CI		
	-3					C									C		C	C	C	C		
XC4002A	-5					CI									CI							
	-5					C									C							
XC4003	-6					CI									CI *		CIM					
	-5					C									C							
XC4003A	-5					1Q93		1Q93	1Q93 *						1Q93							
	-5					1Q93		1Q93	1Q93						1Q93							
XC4004A	-6					1Q93		1Q93							1Q93		1Q93					
	-5					1Q93		1Q93							1Q93		1Q93					
XC4005	-6					CI									CIM	CI	MB				CI	
	-5					C									C	C					C	
XC4005A	-6					1Q93		1Q93							1Q93	1Q93					1Q93	
	-5					1Q93		1Q93							1Q93	1Q93					1Q93	
XC4006	-6														1Q93	1Q93					1Q93	
	-5														1Q93	1Q93					1Q93	
XC4008	-6																			CI	CI	
	-5																			C	C	
XC4010	-6																			CIMB	MB	CI
	-5																			C	C	C

X2824

EPLD Component Availability

	Speed	PC44	WC44	PC68	WC68	PC84	WC84	PG84
XC7236	-30	CI	CI					
	-25	CI	CI					
XC7272	-30				CI	CI	CI	CI
	-25				CI	CI	CI	CI

WC = Windowed

X2841

Symbol	Type	Temperature Range
C	Commercial	0°C - 70°C
I	Industrial	-40°C - 80°C
M	Military	-55°C - 125°C
B	MIL-STD-883, Class B	
*	Top Brazed Ceramic CQFP (CB100)	

XC3100 Switching Characteristics

Description	Symbol	Speed Grade	-5	-4	-3	Units
			Min	Max	Min	Max
Propagation Delays (Input) Pad to Direct In (i) Pad to Registered In (q) with latch transparent Clock (ik) to Registered In (q)	T_{PD} T_{PRG} T_{BR}	3 4	2.8 16.0 2.8		2.5 15.0 2.5	
Set-up Time (Input) Pad to Clock (ik) set-up time	T_{PICK}	15.0		14.0	12.0	
Propagation Delays (Output) Clock (ok) to Pad same (fast) Output (o) to Pad same (slow rate limited) 3-state to Pad begin hi-Z same (fast) 3-state to Pad active and valid same (slow -rate limited)	T_{OKPO} T_{OKPD} T_{OPF} T_{OPS} T_{TSHZ} T_{TSHZ2} T_{TSON} T_{TSOFF}	7 7 10 10 9 9 8 8	5.5 14.0 4.1 13.0 6.9 6.9 12.0 20.0		5.0 12.0 3.7 11.0 6.2 6.2 10.0 17.0	4.4 10.0 3.3 9.0 5.5 5.5 9.0 15.0
Set-up and Hold Times (Output) Output (o) to clock (ok) set-up time Output (o) to clock (ok) hold time	T_{OKO} T_{OKD}	5 6	6.2 0	5.6 0	5.0 0	
Clock Clock High time Clock Low time Max. flip-flop toggle rate	T_{IOH} T_{IOL} F_{CLK}	11 12	2.5 2.5 160	2.2 2.2 195	2.0 2.0 230	
Global Reset Delays (based on XC3042) RESET Pad to Registered In (q) RESET Pad to output pad (fast) (slew-rate limited)	T_{PRH} T_{PRD} T_{PRP}	13 15 15	18.0 24.0 32.0	15.0 20.0 27.0	13.0 17.0 23.0	ns ns ns

New Software Release

General Enhancements

The next release of the XACT development system software will be shipped in the fourth quarter of 1992. Customers currently under a software support agreement should expect to receive their software update by mid December.

Graphic PC Installation Program

PC users will immediately notice the new graphic-based program that simplifies the installation process. The user can select various installation options. If a drive with insufficient disk space is specified, the installation program immediately issues a warning, so that the user can select another drive or partition the software across multiple drives.

The new installation program also supports incremental installation for adding part or package files at a later date.

XDM Improvements

New Family / Part Options: The selection of family and part fields have been tightly integrated in XDM. This also simplifies the Translate and PlaceRoute menus.

Faster Start-up: XDM no longer scans the disk for supported software each time at start-up. Instead, XDM searches the disk the first time and creates a file called proglst.xdm that contains the supported programs found on the first invocation. Future invocations of XDM will start up significantly faster. XDM also has a ScanDisk command to rescan the disk.

Better Color Options: This release contains several new color palettes in XDE, to provide a greater contrast for highlighting elements, and a better balance between foreground and background.

Much Faster XDE Redraw: XDE users on PCs and workstations will notice a more than 10-times improvement in redraw speed. The table below shows XDE redraw times on different platforms for a full XC4010 design, now well under 30 seconds for all platforms except the DN4500.

Platform	v4.30	v4.12
HP 750	3	62
HP 720	4	84
Sparc2	7	128
'486-33	10	190
Apollo 433t	11	219
'386-33	24	423
DN4500	112	796

XDE Redraw Time for a Full XC4010 (in seconds)

XC2000/3000 Enhancements

Several enhancements have been made to the XC2000 / XC3000 software and are included in the DS501 v3.30 release.

Timing-Improvement Router

The Timing-Improvement routing algorithm that was added in APR v3.20 is now accessible via a command-line option. The purpose of Timing Improvement is to reduce delays on nets that have already been routed. By default, APR performs a single timing-improvement pass after the normal routing phase is completed.

Earlier Detection of User Errors

XNFMAP now stops immediately when a user-specified constraint cannot be obeyed. In situations where CLBMAPs, LOC constraints and BLKNM groupings cannot be followed, previous versions of XNFMAP only issued warning messages. These situations now generate errors and give immediate feedback when problems are found.

XC4000 Enhancements

Much of the focus for this release has been on the XC4000 software. Foremost of the new features in the DS502 v1.30 is the XACT-Performance™ feature.

XACT-Performance Software

With the XACT-Performance feature, the user can specify timing requirements in megahertz or nanoseconds on any path in a design, rather than specifying relative weights or criticality on nets alone. The XACT-Performance software partitions, places, and routes designs based on the specified speed requirement. This shortens design cycles by eliminating the iterative method of achieving performance in FPGA designs. Should the user impose an unrealistic requirement on the design, the software responds by issuing a specific error flag indicating the impossible requirement. This helps the designer to focus quickly on the design problem, again shortening the design cycle.

Speed-up with Floating-Point Processor

The PPR program, which is the heart of the XC4000 software, now takes advantage of the 80486 built-in floating-point processor. This cuts design processing time in half. Users with 80387 numeric co-processors will also notice the speed-up in PPR. Due to the added complexity of the XACT-Performance feature, PPR v1.30 does not compile any faster than PPR v1.20 did on '486 machines.

Grouping of FMAPs and HMAPs with Flip-Flops

Users of FMAPs and HMAPs can now employ BLKNM constraints to build complete XC4000 CLBs from logic on their schematic. FMAPs, Flip-Flops and an HMAP with the same BLKNM will be placed in the same CLB. JRD

'486 vs Sparc2 vs HP700

We have run several hundred real XC3000 designs on different platforms. Here is our simple conclusion.

The HP 700 has 2.2 times the speed of the Sparc2, which, in turn, has 2.2 times the speed of a '486-25 PC. In other words, the HP700 beats a '486-25 by a factor of 5.

A '486-50 runs roughly twice as fast as its 25 MHz cousin, which puts the '486-50 performance close to that of a Sparc2. More specifically, the '486-50 is faster in placement, which is the dominant factor for small and simple designs, while the Sparc2 is faster in routing, which is the dominating factor in complex XC3090 designs. The performance on these two platform is within 15% of each other.

VICTORY* for Boundary Scan

The Boundary Scan Description Language (BSDL) files for the XC4000 family are now available and have been verified to be compatible with Teradyne's VICTORY software.

VICTORY, along with other third-party test software, generates test vectors for boundary-scan board-level and device-level testing. One VICTORY module generates patterns for verifying the 1149.1 circuitry and the associated BSDL files. Other modules automatically generate test patterns for Boundary In-circuit, Virtual Interconnect, Virtual Component/Cluster, and Boundary Functional Testing. BSDL is a subset of VHDL, and provides a standard syntax for describing the topology of the test circuitry.

* VICTORY is a trademark of Teradyne Corp.

Legal Protection for Configuration Bit-Stream Programs

The bit-stream program loaded into the LCA may qualify as a "computer program" as defined in Section 101, Title 17 of the United States Code, and as such may be protectable under the copyright law. It may also be protectable as a trade secret if it is identified as such. We suggest that a user wishing to claim copyright and/or trade secret protection in the bit stream program consider taking the following steps.

1. Place an appropriate copyright notice on the LCA device or adjacent to it on the PC board to give notice to third parties of the copyright. For example, because of space limitations, this notice on the LCA device could read "©1992 XYZ Company" or, if on the PC board, could read "Bit Stream © 1992 XYZ Company".
2. File an application to register the copyright claim for the bit-stream program with the U.S. Copyright Office.
3. If practicable, given the size of the PC board, notice should also be given that the user is claiming that the bit-stream program is the user's trade secret. A statement could be added to the PC board such as: "Bit-stream proprietary to XYZ Company. Copying or other use of the bit-

TI and NSC have also used VICTORY to verify their 1149.1-compliant products.

BSDL files for all device / package combinations of the XC4003, XC4005, and XC4010 are available on the Xilinx Technical Bulletin Board (see the Programmable Gate Array Data Book, page 4-4 and 4-5 for access details).

stream program except as expressly authorized by XYZ Company is prohibited."

4. To the extent that documentation, data books, or other literature accompanies the LCA device containing the bit-stream program, appropriate wording should be added to this literature providing third parties with notice of the user's claim of copyright and trade secret in the bit-stream program. For example, this notice could read: "Bit-Stream © 1992 XYZ Company. All rights reserved. The bit-stream program is proprietary to XYZ Company and copying or other use of the bit-stream program except as expressly authorized by XYZ Company is expressly prohibited."
5. To help prove unauthorized copying by a third party, additional non-functional code should be included at the end of the bit-stream program. Therefore, should a third party copy the bit-stream program without proper authorization, if the non-functional code is present in the copy, the copier cannot claim that the bit-stream program was independently developed.

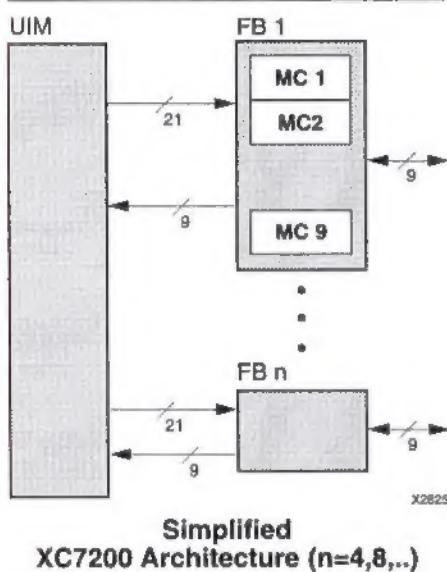
These are only suggestions and Xilinx makes no representations or warranties with respect to the legal effect or consequences of the above suggestions. Each user is advised to consult legal counsel with respect to seeking protection of the user's bit-stream program and to determine the applicability of these suggestions to the user's products and circumstances.

If the user has any questions, contact the Xilinx legal department at 408-879-4984.

Anatomy of the EPLD Architecture

The XC7200 Architecture

The XC7200 devices provide multiple Function Blocks (FBs) interconnected by a central Universal Interconnect Matrix (UIM™). Each FB receives 21 signals from the UIM and produces nine output signals to pins and back into the UIM.

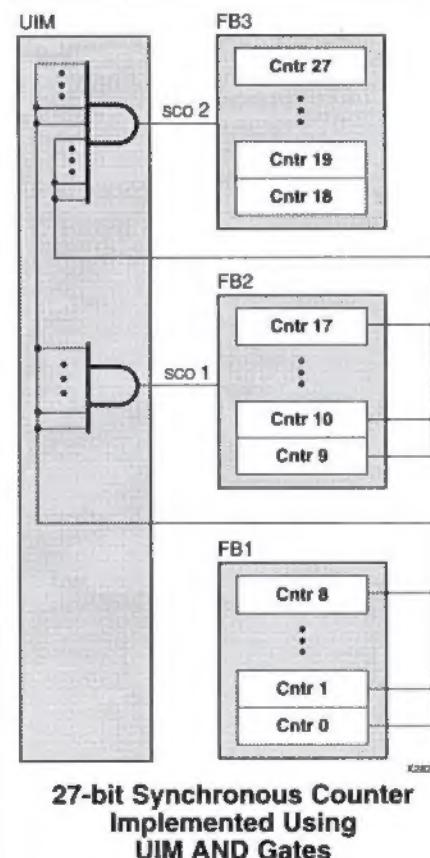


Within each FB there are nine macrocells, each driven by product terms derived from the 21 UIM inputs. Each macrocell resembles a 21V9 PLD architecture. In addition, each macrocell includes an Arithmetic Logic Unit (ALU) that can generate and propagate arithmetic-carry signals between adjacent macrocells and Functional Blocks.

Universal Interconnect Matrix

Unlike other interconnect techniques, Xilinx EPLD's Universal Interconnect Matrix (UIM) provides 100% interconnectivity. Any output of any Function Block can be connected to any input or any number of inputs of any other Function Block using the UIM. The patented interconnect scheme of the UIM provides a fast uniform delay through any of its paths. This interconnect is independent

of fan-in or fan-out loading. Because each FB has identical timing characteristics and the UIM has a constant delay, logic mapped into the device has predictable performance, independent of placement and routing. The UIM can also act as one or more AND gates, e.g., to form terminal count signals within the interconnect. The following diagram illustrates how 27 macrocells can be configured to implement a 27-bit synchronous counter.



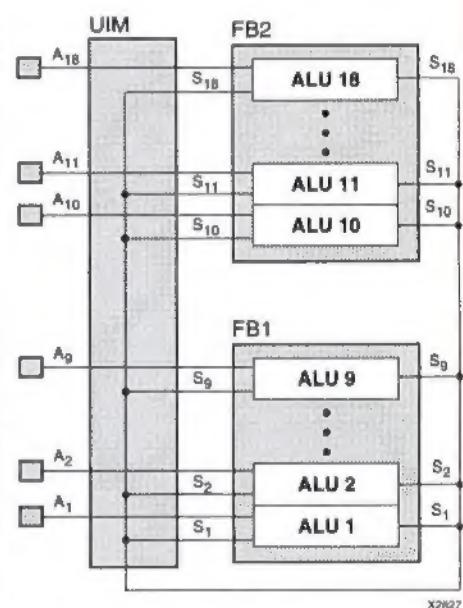
27-bit Synchronous Counter
Implemented Using
UIM AND Gates

Consistent timing performance for all on-chip signals greatly simplifies the design process. In addition to interconnection, the UIM is also used for the following functions:

- Emulating 3-state buses
- Enable/disable signal gates
- Logic decoders
- DeMorgan OR gates

Arithmetic Logic Unit (ALU)

Unlike other programmable logic arrays, the XC7200 architecture includes dedicated arithmetic logic units and fast carry lines running directly between adjacent macrocells and Function Blocks. This additional ALU enables the XC7200 architecture to support fast adders, subtractors, and magnitude comparators of any length up to 72 bits. The following diagram illustrates how 18 macrocells (2 Function blocks) can be programmed to implement an 18-bit accumulator.



18-bit Accumulator
Implemented using ALU Chain

The above architectural features introduce innovative systems-oriented enhancements to the classical features of the PAL-like CPLD architectures. This favors the implementation of fast state machines, large synchronous counters and fast arithmetic, as well as multi-level general-purpose logic.

DR

Dynamic Power Consumption

It is impossible to give a max value for LCA power consumption, because it is totally dynamic. The power consumption of any node is proportional to its capacitance multiplied by the frequency at which it is charged and discharged between +5V and ground. To determine total power, you must know the capacitance of each node and clock line inside the chip, and the frequency with which it is moving up and down; and, you must know the external capacitive load and its frequency.

A worst-case maximum number would be very high, and therefore meaningless, because nobody designs a system where every node moves at 60 MHz, for example.

Estimating power consumption usually has one of two goals: Thermal reliability evaluation, or power-supply sizing.

Thermal calculations can often be substituted by rough estimates, since CMOS power is so low. We give Θ_{JA} values for each package, describing the thermal impedance, i.e. the temperature rise in °C per Watt of power dissipation. Assuming a very conservative max junction temperature of 145°C, a max ambient temperature of 60°C, and a Θ_{JA} of 30°C/W gives a max allowable power dissipation of 2.8 W. Very few LCA designs consume that amount of power, most use a few hundred milliwatts, which results in a junction temperature only a few degrees above ambient.

LCA devices are usually not the dominating power consumers in a system, and do not have a big impact on the power supply design. There are, of course, exceptions to these general rules, and the designer should then use the data on this page to estimate power consumption more accurately. PA

Here are the results of recent measurements of the dynamic power consumption in various Xilinx devices.

Page 6-10 of the 1992 Data Book describes the same parameters, but those values were based on 1988 measurements of devices with larger geometries.

XC2018 at 5.0 V

One CLB driving 3 local interconnects	0.22 mW / MHz
One device output with a 50 pF load	2.0 mW / MHz
One Global clock buffer & line	3.2 mW / MHz

XC2018 at 3.0 V

One CLB driving 3 local interconnects	0.1 mW = 0.03 mA/MHz
One device output with a 50 pF load	0.4 mW = 0.14 mA/MHz
One Global clock buffer & line	1.0 mW = 0.3 mA/MHz

XC3020

One CLB driving 3 local interconnects	0.25 mW / MHz
One device output with a 50 pF load	1.25 mW / MHz
One Global clock buffer & line	2.0 mW / MHz
One Longline <u>without driver</u>	0.1 mW / MHz

XC3090

One CLB driving 3 local interconnects	0.25 mW / MHz
One device output with a 50 pF load	1.25 mW / MHz
One Global clock buffer & line	3.5 mW / MHz
One split Longline <u>without driver</u>	0.15 mW / MHz

XC4003

One CLB driving 3 local interconnects	0.30 mW / MHz
One device output with a 50 pF load	1.2 mW / MHz
One Global clock buffer & line	1.9 mW / MHz
One split Longline <u>without driver</u>	0.12 mW / MHz

XC4005

One CLB driving 3 local interconnects	0.30 mW / MHz
One device output with a 50 pF load	1.2 mW / MHz
One Global clock buffer & line	3.2 mW / MHz
One split Longline <u>without driver</u>	0.17 mW / MHz

XC4010

One CLB driving 3 local interconnects	0.30 mW / MHz
One device output with a 50 pF load	1.2 mW / MHz
One Global clock buffer & line	5.1 mW / MHz
One split Longline <u>without driver</u>	0.24 mW / MHz

Volatility

Xilinx FPGAs use latches to store the data determining logic configuration and interconnects. Configuration information is written into these latches after power has been applied to the device, or whenever a re-configuration is initiated. Obviously, all configuration information is lost if power is interrupted. Some users have voiced concern about this. Here is a detailed explanation.

Configuration information remains valid provided Vcc stays >2.0 V. XC3000 and XC4000 devices, however, have an internal sensor that detects a Vcc drop below a critical value (~3 V). Even though the configuration is valid when that trip point is reached, the device goes into shut-down mode where it 3-states all outputs and clears the configuration memory, preparing it for a reconfiguration when Vcc returns to a more normal value.

There is no possibility of a Vcc dip causing the device to malfunction, i.e., to operate with erroneous configuration information.

- If Vcc stays above the trip point, the device functions normally, albeit at reduced speed, like any other CMOS device.
- If Vcc dips below the trip point, the device 3-states all outputs and waits for reconfiguration.

Some users feel uncomfortable with logic and interconnects defined by the content of latches. There is a concern about accidental or spontaneous changes. Xilinx designers have addressed these concerns. The Xilinx configuration storage latches are simple and rugged, far more rugged than the latches used in typical SRAMs.

Xilinx configuration latches consist of cross-coupled inverters

with active pull-down n-channel and active pull-up p-channel transistors. The High and the Low level are thus both defined with active devices, each having an impedance of ~5 kΩ. Typical SRAMs use passive polysilicon pull-up resistors with an impedance of about 5,000 MΩ. A current of one nanoamp (!) would be sufficient to upset the typical SRAM cell, whereas it would take a million times more current to upset the Xilinx configuration latch.

This does not mean that SRAMs are unreliable, it just shows that the levels in Xilinx configuration latches are six orders of magnitude more resistant to upsets caused by external events, like cosmic rays or alpha particles. Xilinx has never heard about any occurrence of a spontaneous change in the configuration store in any of the 13 million LCA devices sold over the past seven years.

Xilinx production tests the Vcc-dip tolerance of all XC3000 devices in the following way.

- After the device is configured, Vcc is reduced to 3.5 V, and then raised back to 5.0 V. Configuration data is then read back and compared against the original configuration bit stream. Any discrepancy results in rejection of the device.
- Subsequently, Vcc is reduced to 1.5 V and then raised to 5.0 V. The device must first go 3-state, then respond with a request for re-configuration.

Both these tests are performed at high temperature (>70°C for commercial parts, >125°C for military). Any part failing any of these tests is rejected as a functional failure.

PA

Battery Backup

Page 6-18 of the XC3000/2000 Data Book describes a battery-backup circuit that keeps the LCA configured during loss of primary power. Here are some additional comments:

The Seiko voltage detector is called the S-8054 ALB. It comes in a TO-92 package, is available from Seiko Instruments, and is priced at \$ 0.50 to \$ 0.25, depending on quantity. In N. America, call (408) 433-3208; in Europe, call (49) 69 663 0000; in Asia, call (81) 3 3682 5201

Technical Data

Detect Voltage	3.995 V min 4.305 V max
Hysteresis	208 mV typ
Temp. Coeff.	0.52 mV/°C
I _{CC} @ + 6V	2.6 μA typ

Different from the diagram shown in our Data Book, we recommend that pin 2 of the Seiko Voltage Detector be connected to the Vcc pin of the LCA device, not to the incoming primary supply voltage. This insures that the voltage detector device always has a valid supply voltage, and thus can maintain a Low level on the LCA PWRDWN pin.

Seiko also manufactures the S-8420, called the "Battery-Backup IC for 1-Chip CPU", an 8-pin miniDIP device that combines many functions.

- 5-V regulator for up to 50 mA.
- Low-voltage detector (4.2 V)
- Switch-over mechanism
- Three open-drain control outputs.

This looks like an elegant, compact, low-cost solution for small systems with battery backup. We have no experience with this circuit, but will review it in a later issue of XCELL

PA

Xilinx Guarantees Minimum Logic Delays

IC manufacturers traditionally specify and guarantee only the maximum (worst-case) value of delay parameters. Users, however, clamor for double-ended specifications. They argue that assuming zero for the minimum values can lead to theoretical logic hazards and clock-skew-induced hold-time problems. Conscientious designers cannot brush such issues aside, trusting that physical reality will always avoid these theoretical problems.

As a service to the user community, Xilinx now guarantees that the minimum value for all **commercial** speed grades will always be larger than 25% of the max value specified for the fastest documented commercial speed grade.

Example: The shortest combinatorial delay through an XC3000 CLB (T_{ILO}) is 1.4 ns for all commercial speed grades (-50, -70, -100 and -125), since that is 25% of the

5.5 ns max value specified for the fastest speed grade, -125.

For strange historical reasons, set-up times are usually listed in the min column, but the best-case values still follow the same rule.

The shortest set-up time for the logic variables on the a...e inputs is, therefore, 1.5 ns for all XC3000 commercial speed grades. Whenever a non-zero positive hold time is specified for a particular parameter, as it is for DI, the above determination of a shortest set-up time is obviously invalid.

Caution. These minimum delays are guaranteed by design, they are not 100% production-tested.

These minimum values only apply to logic and clock or other buffer delays, not to interconnect delays. The designer should not count on predictable min values for interconnects. Guaranteeing such values would encourage hazardous design practices.

PA

Master & Slave Configure Together

All LCA users should know that daisy-chained devices automatically finish configuration together, and become active simultaneously. Each device counts all CCLK pulses, and each device has its own identical copy of the common length count value. When the number of CCLK pulses received equals this value, all devices in the daisy chain start up together. After a certain number of CCLK pulses, as determined by configuration options, all DONE pins go High, all RESETs are released, and all outputs go active simultaneously. This CCLK-driven synchronous start-up is automatically performed by the configuration control logic.

This information is not new, we only repeat it here because we got some phone calls that showed unnecessary concern on the part of the user.

PA

XC4000 Data Book Reprint

The August 1992 reprint of the XC4000 data book adds pin-out descriptions for the XC4002A, XC4003A, XC4004A, XC4005A, and XC4006. There is a new package outline for the 208 MQFP (Metal Quad Flat Pack). We also added a full page of timing diagrams for the CLB-RAM option (page 45). Beyond that, we corrected a few timing parameters and a small number of typographical errors.

Double-Density PROM

The XC17128 serial configuration PROM offers twice the capacity of an XC1765 at less than twice the price, while retaining the XC1765 functionality, pin-out and also its packaging options: 8-pin DIP and 20-pin PLCC. Speed has been enhanced to 10 MHz, compatible with the fast configuration clock option in the XC4000 family.

Storing max 128 K bits, i.e. 131,072 bits, a single XC17128 can handle the configuration bitstream of one XC4006, or two XC4003s, or two XC3090s, or four XC3042s.

Xilinx Training

We now have more than 20 regional training centers worldwide where we hold classes on a regular basis. You can save travel time and expense by attending a local class. In fact, we can bring any of our classes, even customized for your needs, right into your own facility.

For further information, please contact the local Xilinx sales office, or the Xilinx Training Coordinator at

Tel: (408) 879-5090
Fax: (408) 559-7114

PC-Board Design Hints for LCA Users

Twenty years ago, CMOS was hailed as the friendliest form of logic: no input current, full rail-to-rail logic swing, high noise immunity, soft edges, low power consumption, tolerance for large V_{cc} variations, etc.

Things have changed. CMOS devices have lost some of their user-friendliness as they have become faster and faster, partly in a deliberate quest for speed, partly as an unavoidable result of the smaller device geometries that are required to lower manufacturing cost. The output edge rate is now faster than for TTL, and PC-board interconnect lines between modern CMOS devices can no longer be treated as short circuits or lumped capacitances. The CMOS designer

must now cope with the same transmission-line effects that concerned the previous generation of designers using Schottky-TTL or ECL devices.

Here are some basic rules.

Rule 1: Any PC-board trace is really a transmission line with distributed capacitance and inductance. The series resistance is usually unimportant. The table at left lists typical values for the capacitance and inductance of a PC-board trace with a ground-plane below it.

Any voltage change on such a transmission line causes a corresponding current change. The voltage-to-current ratio is called the characteristic impedance, Z_0 . It is determined by the line thickness and width, by the distance to the ground plane, and by the dielectric constant ϵ of the PC-board material. Z_0 is independent of line length. The table shows typical values for popular constellations.

Rule 2: Signals travel along a transmission line at roughly half the speed of light, or 6" (15 cm) per nanosecond. More precisely, the true propagation speed is the free-air speed of light divided by the square root of the effective dielectric constant, ϵ . The speed of light is very close to 12" (30 cm) per nanosecond, and ϵ for typical epoxy material is 4.7. Since some of the electric field passes through air, the effective ϵ is closer to 4, which leads to the rule of thumb, "half the speed of light".

Rule 3: Whenever the one-way propagation time along a wire or PC trace is longer than half the rise or fall time of the driving signal, this wire or trace must be considered a transmission line, not a lumped capacitive load.

If the rise- or fall-time is 1.5 ns, any PC-board trace longer than 4.5 inches (11 cm) must be analyzed for transmission-line effects.

If the rise or fall time is 5 ns, only PC-board traces longer than 15" (38 cm) need to be analyzed for transmission-line effects.

When a fast rising edge is being driven onto a long transmission line, the driver sees the characteristic impedance Z_0 (50 to 150 Ω), and generates a voltage step that is determined by the ratio of output impedance, R_i , to Z_0 . Typically, an LCA device with an output impedance of 60 Ω drives a 3.5-V step onto a 100- Ω line.

This step propagates to the end of the line at a speed of 6" (15 cm) per nanosecond. If the far end is left open or has a light capacitive load, e.g., the input to a CMOS device, a reflected wave is superimposed on the incoming wave, since only an equal-amplitude reflected wave satisfies the zero-current requirement at the end of the line. This reflected wave travels back to the signal source, arriving there with almost double the original amplitude, usually well above V_{cc}. If the output impedance of the driver differs from Z_0 , the incoming wave is again reflected, travels to the far end, where it is reflected again, etc. This series of reflections with decreasing amplitude is commonly called "ringing". Theoretically, these are rectangular steps of alternating and decreasing amplitude, but high-frequency imperfections often give it the appearance of a decaying sine wave.

At best, such reflections result in poor signal quality and loss of noise immunity. At worst, they reduce system performance and cause functional failures due to double clocking.

0.062" Board			0.031" Board		
Z ₀	Width	C	Width	C	
(Ω)	(mils)	(pF/ft)	(mils)	(pF/ft)	
50	103	35	47	31	
60	77	29	35	27	
70	57	25	26	23	
80	42	22	19	20	
90	31	20	14	18	
100	23	18	10	16	

Imperial units

1.6 mm Board			0.8 mm Board		
Z ₀	Width	C	Width	C	
(Ω)	(mm)	(pF/cm)	(mm)	(pF/cm)	
50	2.6	1.15	1.2	1.02	
60	2.0	0.95	0.9	0.88	
70	1.4	0.82	0.65	0.75	
80	1.1	0.72	0.5	0.66	
90	0.8	0.66	0.35	0.59	
100	0.6	0.6	0.25	0.52	

Metric units

Microstrip-Line Impedance and Capacitance per Unit Length

Coping with Transmission Line Effects

Parallel Termination, Figure 2.

A transmission line of arbitrary length can be terminated at the far end by a resistor to ground or Vcc. If this resistor equals the characteristic impedance Z_0 , the driver always sees the transmission line like a lumped resistive load. Any signal driven onto the line travels to the far end and is dissipated in the resistor. There is no reflection, no ringing or overshoot. Unfortunately, this type of termination is usually impractical, because it puts undue current and power requirements on the driver. It requires 100 mA to drive a 5 V signal onto a 50 Ω line. Only ECL circuits or special buffer circuits can drive terminated transmission lines conveniently. There are two popular methods to alleviate the problem.

- Connecting the terminating resistor through a fairly large capacitor to ground instead of directly to ground or Vcc, reduces static power consumption, but introduces a time constant that must be tailored to the system speed.
- Terminating the line with two resistors, one to ground and one to Vcc, reduces the peak current requirement. 300 Ω to Vcc and 150 Ω to ground is the Thevenin equivalent of a 100 Ω termination to 1.6 V.

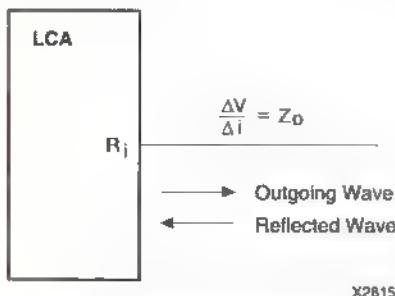


Fig. 1. Transmission Line

Series Termination, Figure 3.

In some cases, series termination at the source can offer the benefits of termination without the drawbacks mentioned above. When an additional series-resistor between the driver and the line increases the effective drive impedance to the same value as Z_0 , the transmission line receives a starting step of half amplitude. Adding an external 40-Ω resistor to the 60-Ω LCA output impedance matches the 100-Ω transmission line, and drives it with a 2.5-V step. This step travels to the far end, where it is reflected and thus doubled in amplitude, as described above. It then arrives back at the driven end of the line with full amplitude (5 V), and is not reflected, since it sees a terminating resistor that is equal to Z_0 .

This seemingly ideal solution has one big drawback: A half-amplitude voltage step travels along the trace and back. Everywhere along the line, except at the far end, this half-amplitude signal can cause trouble, especially in the vicinity of the driver. Series termination is, therefore, recommended only for signals that go from a **single source to a single destination**. Taps on a series-terminated have half-amplitude (2.5 V) levels for fairly long times, which means poor noise immunity and potential malfunction.

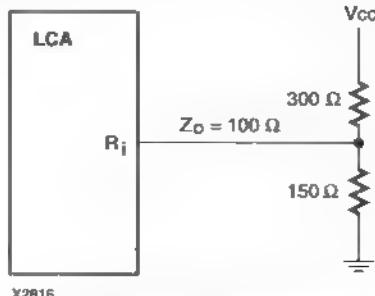


Fig. 2. Parallel Termination

Practical Rules

- Use slew-rate limited outputs wherever possible. Their longer rise and fall times eliminate transmission line effects for all short interconnects.
- Keep critical interconnects as short as possible. It may be better to duplicate some logic in the LCA device and drive from different sides of the device, if that shortens the PC-board traces.
- Use multi-layer PC boards with ground and Vcc planes whenever possible. Always connect all Vcc and ground pins, and be generous with Vcc decoupling capacitors, 0.1 μF per Vcc pin.
- Use series termination for lines that drive a single or lumped destination, but never put taps on a series-terminated line.
- In synchronous systems, the synchronous data and control lines can tolerate poor signal quality **after the clock edge**, but all asynchronous inputs, and especially all clock inputs need good signal quality **all the time**.
- Pay attention to clock distribution on the PC board. Low-skew drivers are now available, e.g., the NSC CGS74C2525.
- CMOS-level input threshold offers the best noise immunity. (Not available on XC4000).
- Remember that a low clock frequency does not make the circuits slow. When the system clock rate is very low, the flip-flops inside the LCA device can still react to 2-ns clock spikes.

PA

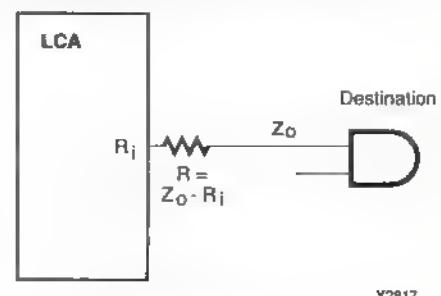


Fig. 3. Series Termination

Crystal-Oscillator Considerations

There are two reasons why many designers feel uncomfortable using the on-chip LCA crystal oscillator circuit.

- This is analog territory, unfamiliar to many digital designers. Words like reactance, transconductance, gain, dB, phase response, $j\omega$ and s-plane evoke memories of long-forgotten early college classes.
- IC documentation is usually skimpy on the issue of specifying crystals and designing reliable oscillator circuits.

Here is additional information.

Let's start with some fundamental facts. There is nothing Xilinx-specific about the oscillator circuit. It's a wide-band inverting amplifier, as used in all popular microcontrollers, like the 8051. When a crystal and some passive components close the feedback path, as described on pages 2-13, 14, 70, and 6-10 of our Data Book, this circuit becomes a reliable and stable clock source.

The path from XTAL2 to XTAL1 inside the LCA device (strangely enough, XTAL2 is the input, XTAL1 is the output) is a single-stage inverting amplifier, which means it has a low-frequency phase response of 180°, increasing by 45° at the 3-dB frequency.

Input impedance is 10-15 pF, input threshold is CMOS, but dc bias must be supplied externally through a megohm resistor from XTAL1 to XTAL2.

Low-frequency gain is about 20, rolling off 3dB at 125 MHz.

Output impedance is between 50 and 100 Ω and the capacitance on the output pin is 10 to 15 pF.

Pulse response is a delay of about 1.5 ns and a rise/fall time of about 1.5 ns.

For stable oscillation,

- the loop gain must be exactly one, i.e., the internal gain must be matched by external attenuation, and
- the phase shift around the loop must be 360° or an integer multiple thereof. The external network must, therefore, provide 180° of phase shift.

A crystal is a piezoelectric mechanical resonator that can be modeled by a very high-Q series LC circuit with a small resistor representing the energy loss. In parallel with this series-resonant circuit is unavoidable parasitic capacitance inside and outside the crystal package, and usually also discrete capacitors on the board.

The impedance as a function of frequency of this whole array starts as a small capacitor at low frequencies (Figure 1). As the frequency increases, this capacitive reactance decreases rapidly, until it reaches zero at the series resonant frequency.

At slightly higher frequencies, the reactance is inductive, starting with a zero at series resonance, and increasing very rapidly with frequency. It reaches infinity when the effective inductive impedance of the series LC circuit equals the reactance of the parallel capacitor. The parallel resonance frequency is a fraction of a percent above the series-resonance frequency.

Over this very narrow frequency range between series and parallel resonance, the crystal impedance is inductive and changes all the way from zero to infinity. The energy loss represented by the series resistor prevents the impedance from actually reaching zero and infinity, but it comes very close.

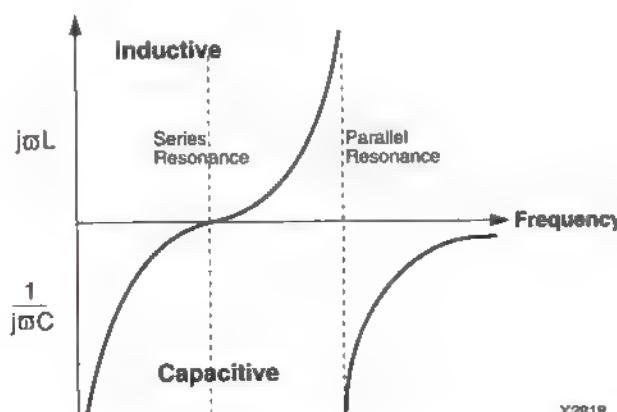
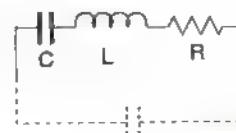


Figure 1. Reactance as a Function of Frequency



X2835

Figure 2. Equivalent Circuit

Microprocessor- and FPGA-based crystal oscillators all operate in this narrow frequency band, where the crystal impedance can be any inductive value. The circuit oscillates at a frequency where the attenuation in the external circuit equals the gain in the LCA device, and where the total phase shift, internal plus external, equals 360°.

Figure 3 explains the function. At the frequency of oscillation, the series-resonant circuit is effectively an inductor, and the two capacitors act as a capacitive voltage divider, with the center-point grounded. This puts a virtual ground somewhere along the inductor and causes the non-driven end of the crystal to be 180° out of phase with the driven end, which is the external phase shift required for oscillation. This circuit is commonly known as a Pierce oscillator.

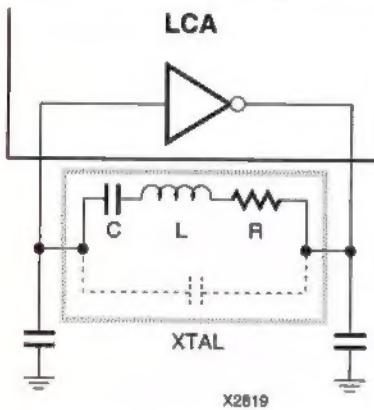


Figure 3. Pierce Oscillator

Series Resonant or Parallel Resonant?

Crystal manufacturers label some crystals as series-resonant, others as parallel-resonant, but there really is no difference between these two types of crystals, they all operate in the same way. Every crystal has a series resonance, where the impedance of the crystal is extremely low, much lower than at any other frequency. At a slightly higher frequency, the crystal is inductive and in parallel resonance with the unavoidable

Practical Considerations

- The series resonance resistor is a critical parameter. To assure reliable operation with worst-case crystals, the user should experiment with a discrete series resistor roughly equal to the max internal resistance specified by the crystal vendor. If the circuit tolerates this additional loss, it should operate reliably with a worst-case crystal without the additional resistor.
- The two capacitors affect the frequency of oscillation and the start-up conditions. The series connection of the two capacitors is the effective capacitive load seen by the crystal, usually specified by the crystal vendor.
- The two capacitors also determine the minimum gain required for oscillation. If the capacitors are too small, more gain is needed, and the oscillator may be unstable. If the capacitors are too large, oscillation is stable but the required gain may again be higher. There is an optimum capacitor value, where oscillation is stable, and the required gain is at a minimum. For most crystals, this capacitive load is around 20 pF, i.e., each of the two capacitors should be around 40 pF.
- Crystal dissipation is usually around 1 mW, and thus of no concern. Beware of crystals with "drive-level dependence" of the series resistor. They may not start up. Proper drive level can be checked by varying Vcc. The frequency should increase slightly with an increase in Vcc. A decreasing frequency or unstable amplitude indicate an over-driven crystal. Excessive swing at the XTAL2 input results in clipping near Vcc and ground. An additional 1 to 2 kΩ series resistor at the XTAL1 output usually cures that distortion problem. It increases the amplifier output impedance and assures additional phase margin, but results in slower start-up.
- Be especially careful when designing an oscillator that must operate near the specified max frequency. The circuit needs excess gain at small signal amplitudes to supply enough energy into the crystal for rapid start-up. High-frequency gain may be marginal, and start-up may be impaired.
- Keep the whole oscillator circuit physically as compact as possible, and provide a single ground connection. Grounding the crystal can is not mandatory but may improve stability.

Sources

Fick: "Schwingquarz und Mikroprozessor". *Elektronik*, Feb. 1987

Horowitz & Hill: *The Art of Electronics*, Cambridge University Press, 1989

Motorola High-Speed CMOS Logic Data Book., 1983.

PA

stray capacitance or the deliberate capacitance between its pins.

The only difference between the two types of crystal is the manufacturer's choice of specifying either of the two frequencies. If series resonance is specified, the actual frequency of oscillation is a little higher than the specified value. If parallel resonance is specified, the frequency of oscillation is a little lower. In most cases, these small deviations are irrelevant.

Simple RC Oscillator

This simple RC oscillator uses the XTAL2 or TCLK input, both of which guarantee a CMOS input threshold. The two counter-phase outputs are driven in such a way that the inverting output is derived from the non-inverting output. This prevents any possibility of spurious oscillations or erroneous operation.

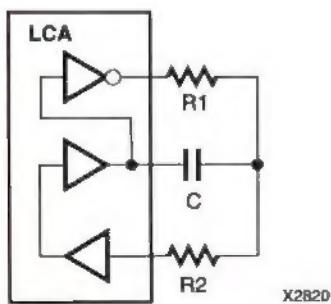
R1 is the main timing resistor. R2 allows the timing node to swing below ground and above Vcc.

$$f = 1 / 2RC \quad R=R_1=R_2$$

The resistors can be any value between $500\ \Omega$ and $1\ M\Omega$. At the low end, the output current becomes excessive; at the high end, the input leakage current affects the frequency, although this leakage current is really much lower than the $10-\mu A$ specification.

The capacitor can have any value between $50\ pF$ and $1\ \mu F$. At the low end, the input capacitance affects the frequency; at the high end, the capacitor may become too big and too expensive. It is better to use an internal binary divider to extend the frequency to a lower value.

PA



X2820

Metastability

Pages 6-16 and 17 of the Xilinx Data Book describe the metastability phenomenon, and show the excellent behavior of Xilinx FPGA internal flip-flops. Some users prefer an equation over the simple graph. Here it is.

The mean time between failures (MTBF) is a statistically defined value. It is inversely proportional to the product of the two frequencies involved, the clock frequency, and the average frequency of the data changes, provided that these two frequencies are independent.

K_1 is a factor that relates to the likelihood of going metastable. The second factor affecting MTBF is an exponential function of the additional delay that must be allowed for settling out of the metastable condition.

Compared to industry-standard Schottky-TTL flip-flops, the Xilinx internal flip flops are a million times less likely to go metastable, and they return to their stable state much faster. An addi-

The Fairchild FAST seminar of 1982 showed the following equations for industry-standard Schottky-TTL devices :

$$\text{for 74F74, MTBF (in seconds)} = (1 / f_1 \times f_2 \times 5.5 \times 10^{-5}) \times e^{1.25 \times t}$$

$$\text{for 74S74, MTBF (in seconds)} = (1 / f_1 \times f_2 \times 20) \times e^{1.25 \times t}$$

Assuming the same conditions (1MHz and 10 MHz) as for the Xilinx devices gives the following values for MTBF:

74F74

MTBF = 50 μs for 10 ns delay

MTBF = 600 μs for 12 ns delay

MTBF = 25 ms for 15 ns extra delay

MTBF = 13 s for 20 ns extra delay

74S74

MTBF = 36 μs for 20 ns extra delay

MTBF = 10 s for 30 ns extra delay

MTBF = 30 days for 40 ns delay

tional delay of 10 to 15 ns is more than adequate to make metastability problems statistically insignificant.

PA

$$\text{MTBF} = (1 / f_1 \times f_2 \times K_1) \times e^{K_2 \times t}$$

Xilinx measurements established: $K_1 = 1.5 \times 10^{-10}$, $K_2 = \ln 40 = 3.69$

$$\text{MTBF} = (1 / f_1 \times f_2 \times 1.5 \times 10^{-10}) \times e^{3.69 \times t} \text{ for Xilinx LCA}$$

MTBF in seconds, f1 and f2 in Hz, t, the additional delay, in ns

$K_1 = 1.5 \times 10^{-10}$ is a constant with the dimension of time

$K_2 = 3.69 = \ln 40$, i.e., the increase by a factor 40 for each additional ns.

- 1MHz, 10 MHz and an additional acceptable delay of 1.0 ns:

$$\text{MTBF} = (1 / 1.5) \times 10^{-3} \times e^{3.69} = 0.67 \times 10^{-3} \times 40 = 27 \text{ milliseconds}$$
- 1MHz, 10 MHz and 4.2 ns:

$$\text{MTBF} = (1 / 1.5) \times 10^{-3} \times e^{15.5} = 0.67 \times 10^{-3} \times 5.38 \times 10^6 = 3600 \text{ s} = 1 \text{ hr}$$
- 1MHz, 10 MHz and 6.7 ns: MTBF = 423 days
- 1MHz, 10 MHz and 8.5 ns: MTBF = 889 years
- 1MHz, 10 MHz and 10.0 ns: MTBF = 225,000 years
- 1MHz, 10 MHz and 11.0 ns: MTBF = 9 million years
- 1MHz, 10 MHz and 12.0 ns: MTBF = 360 million years. Enough?

Shorter Power-on Delay

The technique described below reduces the power-on initialization delay by 40 to 100 ms. Only a few users are concerned about this delay, but for them, this technique offers a simple solution.

All LCA devices check their M0 mode inputs during power-up. Master devices then lengthen the power-on time delay by a factor of four, so that the master always takes longer to get ready than any slave, even a slow slave. This covers worst-case processing variations that might change the internal clock frequency over a 3:1

range. The daisy-chain can thus be configured "open-loop"; the master need not check the readiness of the slave.

XC3000 and XC4000 devices, however, each have an INIT output that can be used to inform the master about the state of readiness of the slaves. Just wire all slave INIT outputs together (excluding the master), and drive the RESET input of the XC3000 master, or the bidirectional INIT input pin of the XC4000 master.

When slaves communicate in this way with the master, there

really is no need for the extra power-on delay built into the master. In an XC4000 master, this delay can easily be shortened to the same value as in a slave. For this purpose, the M0 mode input (or all three mode inputs together) must be driven by the inverse of the wire-ANDED INIT outputs from all slaves, as shown in Fig. 1.

During power-up, the lead device sees a High on M0, and considers itself a slave, which means a short power-on delay (nominally only 22 ms instead of 88 ms). After INIT has gone High, the lead device checks all three mode pins, and now finds M0 Low, causing the lead device to act as a master, activating its CCLK output and proceeding with the configuration process.

This "trick" works with any XC4000 device intended as a serial or parallel master, with or without slaves. Those slaves can be either XC4000 or XC3000 devices, but not XC2000s, since they do not have INIT outputs.

The circuit must be modified for an XC3000 master, since the XC3000 device lacks the microsecond delay between the end of INIT and the sampling of the mode pins. Figure 2 shows an external delay and Schmitt-trigger circuit that achieves the same result.

RESET is held Low for an additional μ s after INIT has gone High. This assures that M0 is Low when it is tested before the start of configuration.

PA

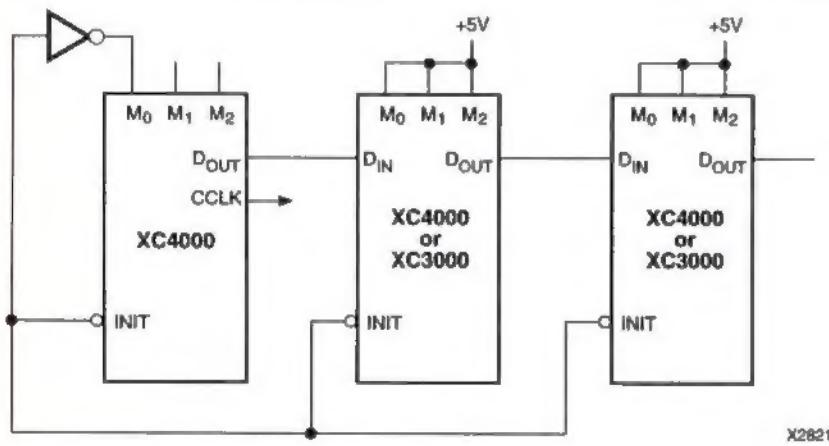


Figure 1. XC4000 master

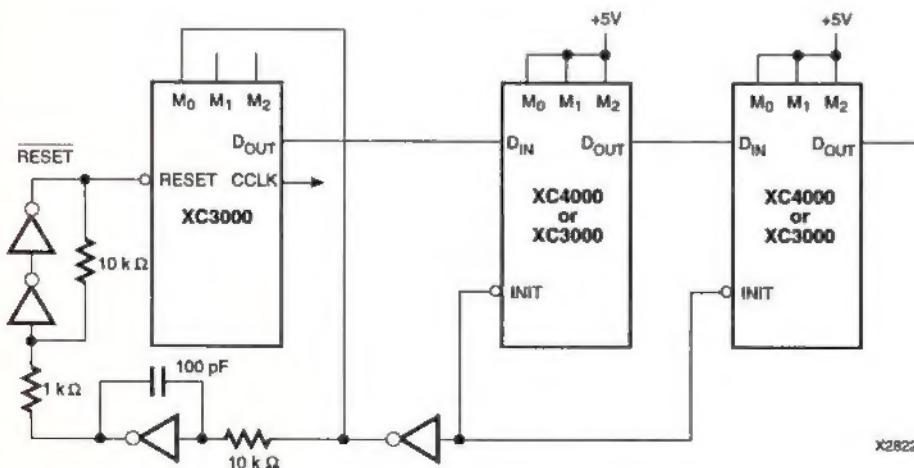


Figure 2. XC3000 master

For the Guinness Book of Records: 255-MHz Frequency Synthesizer, the Fastest FPGA Application Ever

At the Electronica show, Nov. 10 through 14 in Munich, Germany, Xilinx is demonstrating the fastest FPGA application ever. A 1-to-255-MHz crystal-controlled pulse generator is implemented in one quarter of the new XC3130 device. The rest of the FPGA contains the pushbutton-operated control logic and the LCD readout drivers. The only external components are a voltage-controlled oscillator (VCO), built of discrete devices, a 4-MHz reference crystal, a 30-V integrating op-amp, and, of course, the LCD display.

Inside the FPGA, a presetable counter with a divide range of 128 to 255, and an output multiplexer that extends the frequency range down to 1 MHz, are driven by the VCO clock of 128 to 255 MHz. These 15 CLBs are laid out in a speed-optimized pattern.

The control circuitry, including the 1-MHz phase/frequency comparator, is fairly uncritical. The frequency being generated is displayed in BCD seven-segment form on a 0.5" 3-digit liquid crystal display, driven directly by the LCA. Two buttons, up and down, scroll

the frequency at a rate that increases first 4 times, then 16 times, then 32 times, as either button is held down for several seconds.

This design demonstrates the high performance of the newest XC3100 family. 255-MHz operation must be considered the "high water mark" for FPGA designs. Typical performance in synchronous designs is usually in the 50-to-80 MHz system clock range, but small, well-defined subsystems can operate much faster, as shown in this frequency synthesizer application.

PA



2100 Logic Drive
San Jose, CA 95124-3450

U.S. POSTAGE
PAID
FIRST CLASS
PERMIT NO. 2196
SAN JOSE, CA